



UNITED STATES PATENT AND TRADEMARK OFFICE

A

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/608,354	06/29/2000	Donald Hooper	10559/222001/P8715	8914

20985 7590 11/17/2005

FISH & RICHARDSON, PC
P.O. BOX 1022
MINNEAPOLIS, MN 55440-1022

EXAMINER

REILLY, SEAN M

ART UNIT PAPER NUMBER

2153

DATE MAILED: 11/17/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

DETAILED ACTION

This office action is in response to Applicant's amendment and request for reconsideration filed on August 25, 2005. Claims 23-35 are presented for further examination. A new grounds of rejection is set forth in this action, accordingly this action is made NON-FINAL.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

1. Claims 23-35 are rejected under 35 U.S.C. 103(a) as being unpatentable over Lipman et al. (U.S. patent No. 6,192,051, hereinafter "Lipman"), in view of Sindhu et al. (U.S. Patent No. 6,493,347, hereinafter "Sindhu") and Brown (U.S. Patent Number 6,691,218).

In considering claim 23, Lipman discloses network router ("router 10") comprising:

An input switch and an output switch (inherent in a router);

A controller ("controller 12," Fig. 1; col. 6, lines 34-35), the controller comprising a look-up engine receiving look-up requests (col. 7, lines 17-50);

a memory for storing data for access by a longest prefix match program (Fig. 2, "forwarding controller memory"), the program comprising:

A data structure stored in the memory, the data structure including information resident in a database (“routing database”; col. 7, lines 21-23) used by the longest prefix match program and including:

A large table at a root (“Level 1 routing entries”; col. 10, line 34), the root branching to nodes containing small trie tables (“level-2 routing entries”; col. 10, lines 61-62; Fig. 7), each trie table addressed by a span of IP address bits (“IP address bits [15:8]”) to locate an indexed trie entry (col. 10, lines 39-46, 61-63), the indexed trie entry including a route pointer (“next hop address”) and a trie pointer (“key”; col. 10, lines 39-46; col. 11, lines 5-14, wherein the next hop address points to a destination address and the key points to another trie table within the routing structure).

However, Lipman does not teach that the controller has a plurality of look-up engines, each engine receiving look-up requests in a round-robin fashion. Nonetheless, using a plurality of look-up engines in a router to receive routing requests in a round-robin fashion is well-known, as evidenced by Sindhu. In a similar art, Sindhu discloses a router for switching packets from a source to a destination, including a controller wherein “a plurality of route look-up engines 110 are included in controller 106, each receiving look-up requests in round-robin fashion” (col. 16, lines 52-56). Sindhu further describes a motivation for using a plurality of route look-up engines as “to speed the routing process” (col. 16, lines 56-57). Therefore, given this well-known feature and its benefits both disclosed by Sindhu, it would have been obvious to include it in the controller disclosed by Lipman.

Lipman also fails to specifically recite that the look-up engines traverse in parallel the trie tables. Nonetheless, it was well known in the art at the time invention to search trie tables in

Art Unit: 2153

parallel within the routing environment, as evidenced by Brown. In an analogous art, Brown disclosed a routing look-up system that uses multiple trie look-up tables to compute the next hop a packet should be routed to (Col 11, line 61 – Col 12, line 25), similar to the Lipman system. In Brown's system the trie tables are searched in parallel for a given search key (Col 12, lines 42-59). By searching the tables in parallel search time is reduced, which in turn reduces the overall packet routing time. Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention to incorporate the parallel processing disclosed by Brown within Lipman's system, in order to reduce packet routing latency.

In considering claim 24, Lipman further discloses that the small trie tables each comprise prefix match fields for each indexed entry (i.e. "pointer values"; col. 11, lines 46-47), a population count of pointers (i.e. 256 level-2 entries, see Fig. 7), and hidden prefix entries (col. 12, lines 40-50, describing pointers that only act as prefix entries when another pointer is deleted).

In considering claim 25, Lipman further discloses that the hidden prefix entries hold shorter prefix entry pointers (col. 12, lines 51-58, describing the shorter entry pointers being used when the longer pointer is deleted).

In considering claim 26, Lipman further discloses that the small trie tables are stored in SRAM (i.e. cache) and used for route lookups ("lookup"; col. 9, lines 49-51), route adds ("adds"), and route deletes ("deletes"; col. 12, lines 1-14, 40-50).

In considering claim 27, Lipman further discloses that the indexed trie entry is a 32-bit longword (col. 10, lines 54-56; Fig. 7, box 140, showing the 32-bit IP address).

In considering claim 28, Lipman discloses a network router ("router 10") comprising:

A plurality of input and output ports (col. 14, lines 23-34, describing "input" and "output" "ports P0-P3") linked to an input switch and an output switch (both inherent);

A controller ("controller 12," Fig. 1; col. 6, lines 34-35), the controller comprising a look-up engine receiving look-up requests (col. 7, lines 17-50); and

A memory (Fig. 2, "forwarding controller memory"), the memory including a method of searching a database for a prefix representing a destination address (col. 7, lines 21-41, "routing database... enable[s] the device 10 to make decisions regarding how packets received on a segment 20 or 22 are to be forwarded), the method comprising:

Reading a data structure stored in the memory, the data structure comprising a large table at a root ("level-1 routing entries"; col. 10, lines 55-57), the root branching to two nodes containing small trie tables ("level-2 routing entries"; col. 10, lines 61-62; Fig. 7), each trie table addressed by a span of IP address bits ("IP address bits [15:8]") to locate an indexed trie entry (col. 10, lines 39-46, 61-63), the indexed trie entry including a route pointer ("next hop address") and a trie pointer ("key"; col. 10, lines 39-46; col. 11, lines 5-14, wherein the next hop address points to a destination address and the key points to another trie table within the routing structure); and

However, Lipman does not teach that the controller has a plurality of look-up engines, each engine receiving look-up requests in a round-robin fashion. Nonetheless, using a plurality of look-up engines in a router to receive routing requests in a round-robin fashion is well-known, as evidenced by Sindhu. In a similar art, Sindhu discloses a router for switching packets from a source to a destination, including a controller wherein “a plurality of route look-up engines 110 are included in controller 106, each receiving look-up requests in round-robin fashion” (col. 16, lines 52-56). Sindhu further describes a motivation for using a plurality of route look-up engines as “to speed the routing process” (col. 16, lines 56-57). Therefore, given this well-known feature and its benefits both disclosed by Sindhu, it would have been obvious to include it in the controller disclosed by Lipman.

Lipman also fails to specifically recite that the look-up engines traverse in parallel the trie tables. Nonetheless, it was well known in the art at the time invention to search trie tables in parallel within the routing environment, as evidenced by Brown. In an analogous art, Brown disclosed a routing look-up system that uses multiple trie look-up tables to compute the next hop a packet should be routed to (Col 11, line 61 – Col 12, line 25), similar to the Lipman system. In Brown’s system the trie tables are searched in parallel for a given search key (Col 12, lines 42-59). By searching the tables in parallel search time is reduced, which in turn reduces the overall packet routing time. Thus, it would have been obvious to one of ordinary skill in the art at the time of the invention to incorporate the parallel processing disclosed by Brown within Lipman’s system, in order to reduce packet routing latency.

In considering claim 29, Lipman further discloses that the route pointer represents the destination address (“next hop address”) and the trie pointer points to a next small trie table (“key”; col. 10, lines 39-46; col. 11, lines 5-14, wherein the next hop address points to a destination address and the key points to another trie table within the routing structure).

In considering claim 30, Lipman further discloses that the small trie tables each comprise prefix match fields for each indexed entry (i.e. “pointer values”; col. 11, lines 46-47), a population count of pointers (i.e. 256 level-2 entries, see Fig. 7), and hidden prefix entries (col. 12, lines 40-50, describing pointers that only act as prefix entries when another pointer is deleted).

In considering claim 31, Lipman further discloses reporting a non-match if the prefix does not match an entry (col. 12, lines 1-14, “if no such tree or trees exist, then new level-3 and/or level-2 trees are created for the new routing entry”).

In considering claim 32, Lipman further discloses that a first large table is a single 64K entry table that is indexed by bits 31:16 of an IP address (“64k level 1 entries,” “indexed by IP address bits [31:16]”; col. 10, lines 56-58).

In considering claim 33, Lipman does not explicitly disclose that a second large table is indexed specifically by bits 31:24 of an IP address. However, Lipman does disclose that a second large table could be indexed by bits [31:20] and further suggests that “in alternative

Art Unit: 2153

embodiments it may be desirable to shuffle the address fields with respect to the levels” (col. 18, line 65 – col. 19, line 5). Thus, as evidenced by Lipman, the specific selection of bits for indexing each level is a matter of design choice, and it would have been obvious to include only bits 31:24 in the root table taught by the combined system of Lipman and Sindhu (rather than bits 31:16 or 31:20) because fewer indexing bits at any particular level allows traversal through that level to either be faster or to use less memory.

In considering claims 34 and 35, Lipman further discloses that the small tables are dynamically allocated (i.e. added to or deleted from in real time, col. 12, lines 1-50) and comprise:

A tree with each node representing 16 bits of addresses covering an extension of 1-16 bits of a prefix entry from a previous tree (Fig. 7, node 40, for example). Again, although Lipman does not explicitly disclose that the nodes in the small tables represent 4 bits of address covering an extension of 1-4 bits of a prefix entry from a previous tree, Lipman discloses that it may be desirable to shuffle the address fields with respect to the levels. Here, it would have been desirable to shuffle the address fields in level one such that each node at the level-2 table covers only 4 bits of addresses, in order to simplify the level-2 and level-3 (or to even eliminate the need for the level-3) tables. Therefore, it would have been obvious for the small tables taught by the combined system of Lipman and Sindhu to have nodes representing only 4 bits of addresses, instead of 16 bits.

Response to Arguments

Applicant's arguments are noted however they are moot in view of the grounds of rejection set forth.

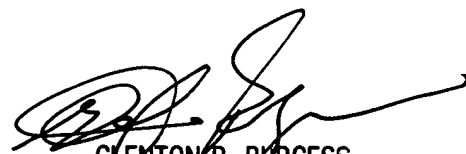
Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Sean Reilly whose telephone number is 571-272-4228. The examiner can normally be reached on M-F 8-5.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Glen Burgess can be reached on 571-272-3949. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

11/9/2005


GLENTON B. BURGESS
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100